Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE R&D DÉFENSE

Maritime Anomaly Detection: Domain Introduction and Review of Selected Literature

Etienne Martineau
Jean Roy
DRDC Valcartier

Defence R&D Canada – Valcartier
Technical Memorandum
DRDC Valcartier TM 2010-460
October 2011

Canada

Maritime Anomaly Detection: Domain Introduction and Review of Selected Literature

Etienne Martineau
Jean Roy
DRDC Valcartier

Defence R&D Canada – Valcartier

Technical Memorandum
DRDC Valcartier TM 2010-460
October 2011

Principal Author

Original signed by Etienne Martineau

Etienne Martineau

Defence Scientist

Approved by

Original signed by Stéphane Paradis

Stéphane Paradis

Section Head / Intelligence & Information Section, DRDC Valcartier

Approved for release by

Original signed by Christian Carrier

Christian Carrier

Chief Scientist, DRDC Valcartier

DRDC project 11hg

Abstract

Early in the conduct of Project 11hg (Collaborative Knowledge Exploitation for Maritime Domain Awareness) at Defence R&D Canada, anomaly detection in the maritime domain was identified by the operators/analysts of the operational community as an important aspect requiring research and development. A number of R&D activities have thus been undertaken under the project to specifically investigate maritime anomaly detection (MAD). This Technical Memorandum reports on one of these activities. It first provides a high-level introduction to the domain, and then presents a review of selected literature on the subject. Different views of the field are presented, starting with a description of the various steps of MAD, followed by a discussion of four interrelated goals of MAD. Current gaps in MAD are identified from the data and information, processing and systems perspectives. The selected literature review is structured around specific organizations known to be active in maritime anomaly detection, various MAD systems, and other relevant research activities. A high-level assessment of the methods for MAD that were found in the reviewed literature completes the discussion.

Résumé

Tôt dans l'exécution du projet 11hg (Collaborative Knowledge Exploitation for Maritime Domain Awareness) à Recherche et développement pour la défense Canada, la détection d'anomalies dans le domaine maritime a été identifiée par les opérateurs/analystes de la communauté opérationnelle comme un aspect important qui nécessite de la recherche et du développement. Plusieurs activités de R et D ont été entreprises dans le cadre du projet pour étudier spécifiquement la détection d'anomalies maritimes (DAM). Ce mémorandum technique fait état de l'une de ces activités. Il fournit d'abord une introduction au domaine et présente ensuite une revue de la littérature sélectionnée sur le sujet. Différentes visions du domaine sont présentées, en commençant par la description des différentes étapes de la DAM suivie d'une discussion sur ses quatre objectifs. Les lacunes actuelles de la DAM sont identifiées dans la perspective des données et de l'information, du traitement et des systèmes. La revue de la littérature est structurée autour d'organisations spécifiques reconnues comme actives dans la DAM, des divers systèmes de gestion de la DAM et d'autres activités pertinentes de recherche. Une évaluation de haut niveau des méthodes de la DAM qui ont été identifiées dans la revue de la littérature complète la discussion.

This page intentionally left blank.

Executive summary

Maritime Anomaly Detection: Domain Introduction and Review of Selected Literature

Etienne Martineau; Jean Roy; DRDC Valcartier TM 2010-460; Defence R&D Canada – Valcartier; October 2011.

This Technical Memorandum documents results and findings from research activities conducted at Defence R&D Canada (DRDC) under Project 11hg. This project, entitled "Collaborative Knowledge Exploitation for Maritime Domain Awareness," is part of DRDC's Applied Research Program (ARP). Its objective is to explore and develop an integrated collaborative knowledge management and exploitation (KME) technology framework to allow operators/analysts to quickly find, access, create, organize, share, use and reuse relevant knowledge for maritime domain awareness in the Regional Joint Operations Centres (RJOCs) on the East and West coasts of Canada, and the Joint Command Centre (JCC) of Canada Command (CANCOM) Headquarters. Expected outcomes for the project are advanced KME capabilities supporting the staff in the building of maritime situational awareness through knowledge discovery, automated reasoning, and situation analysis and assessment.

Early in the conduct of Project 11hg, anomaly detection in the maritime domain was identified by the operators/analysts as an important aspect requiring R&D. Critical mandates such as defending sovereignty, protecting infrastructure, countering terrorism and detecting illegal activities have all become more challenging in the maritime domain. A large portion of the information made available to the staff originates from platforms going about normal, legitimate activities, and identifying anomalous events worthy of attention is a very demanding task. Given the importance of anomaly detection for the operational community, R&D activities have been undertaken under Project 11hg to investigate the subject.

This Technical Memorandum reports on one of these R&D activities. It first provides a high-level introduction to the domain of maritime anomaly detection (MAD), then presents a review of selected literature on the subject. Different views of the domain are presented, starting with a description of the various steps of MAD, followed by a discussion of four interrelated goals of MAD. Current gaps in MAD are identified from the data and information, processing and system perspectives. The selected literature review is structured around specific organizations known to be active in maritime anomaly detection, various MAD systems, and other relevant research activities. A high-level assessment of the methods for MAD that were found in the reviewed literature completes the discussion.

Most of the R&D activities on MAD conducted under Project 11hg revolve around the use of knowledge-based system technologies. Driven by the objective of exploring the other avenues being pursued by the community active in MAD or related matters, the work reported here is thus important as it is somewhat complementary to the current efforts in Project 11hg.

In light of the findings documented here, it is expected that some of the techniques and methods for MAD will be further investigated and eventually developed to be integrated with the proof-of-concept prototypes already implemented under Project 11hg.

Sommaire

Maritime Anomaly Detection: Domain Introduction and Review of Selected Literature

Etienne Martineau; Jean Roy; DRDC Valcartier TM 2010-460; R & D pour la défense Canada – Valcartier; Octobre 2011.

Ce mémorandum technique documente les résultats et les découvertes des activités de recherche effectuées à Recherche et développement pour la défense Canada (RDDC) dans le cadre du projet 11hg. Ce projet, intitulé « Collaborative Knowledge Exploitation for Maritime Domain Awareness », fait partie du Programme de recherche appliquée (PRA). Son objectif est d'explorer et d'intégrer des infrastructures technologiques de gestion et d'exploitation de la connaissance (GEC) pour permettre rapidement aux opérateurs/analystes de trouver, accéder à, créer, organiser, partager et utiliser les connaissances pertinentes pour l'éveil situationnel maritime dans les centres régionaux d'opérations interarmées (CROI) et au Centre de commandement interarmées (CCI) du commandement canadien (CANCOM). Les résultats attendus pour le projet sont des avancées dans les capacités de GEC qui supportent le personnel dans l'éveil situationnel par la découverte de connaissances, le raisonnement automatisé, l'analyse et l'évaluation de la situation.

Tôt dans l'exécution du projet 11hg, la détection d'anomalies dans le domaine maritime (DAM) a été identifiée par les opérateurs/analystes comme un aspect important qui nécessite de la R et D. Les mandats critiques, tels que défendre la souveraineté, protéger les infrastructures, combattre le terrorisme et détecter les activités illégales sont maintenant plus exigeants dans le domaine maritime. Une grande portion des informations fournies au personnel provient de plateformes exécutant des activités normales et légitimes et il est très exigeant d'identifier les événements dignes d'attention. Étant donné l'importance de la DAM pour la communauté opérationnelle, des activités de R et D ont été entreprises dans le cadre du projet 11hg pour étudier le sujet.

Ce mémorandum technique fait état de l'une de ces activités de R et D. Il fournit d'abord une introduction au domaine de la DAM et présente ensuite une revue de la littérature sélectionnée sur le sujet. Différentes visions du domaine sont présentées, en commençant par la description des différentes étapes de la DAM suivie d'une discussion sur quatre objectifs de la DAM. Les lacunes actuelles de la DAM sont identifiées dans la perspective des données et de l'information, du traitement et des systèmes. La revue de la littérature est structurée autour d'organisations spécifiques reconnues comme actives dans la DAM, des divers systèmes de gestion de la DAM et d'autres activités pertinentes de recherches. Une évaluation de haut niveau des méthodes de la DAM qui ont été identifiées dans la revue de la littérature complète la discussion.

La majorité des activités de R et D effectuées dans le cadre du projet 11hg tournent autour de l'usage de systèmes à base de connaissances. Motivé par la volonté d'explorer d'autres avenues que celles suivies par la communauté active dans la DAM ou dans des thèmes apparentés, le travail présenté ici est donc important, car il est complémentaire aux efforts dans le projet 11hg.

À la lumière des découvertes documentées ici, il est souhaitable que quelques méthodes ou techniques de la DAM soient plus approfondies et qu'elles soient finalement intégrées aux prototypes de validation de principe déjà développés dans le cadre du projet 11hg.

Table of contents

List of figures

List of tables

This page intentionally left blank.

1. Introduction

Early in the conduct of Project 11hg at Defence R&D Canada, anomaly detection in the maritime domain was identified by the operators/analysts of the operational community as an important aspect requiring research and development. A number of R&D activities have therefore been undertaken under the project to specifically investigate maritime anomaly detection (MAD). The general objective of this document is to familiarize the reader with the current status of maritime anomaly detection. This document is not state of the art; it only provides some material necessary for understanding the current issues and requirements of the field. In addition, recent studies are summarized in Part 3 in order to make readers aware of recent research efforts.

In Part 2, our goal is to capture the essence of recent presentations and workshops on maritime anomaly detection (MAD). Part 2 is not an exhaustive literature review but rather a summary of current topics being discussed by researchers and in the industry and thus provides a high-level introduction to the domain. Part 3 is a review of selected papers addressing the issues raised in Part 2. Different views of the field are presented, starting with a description of the various steps of MAD, followed by a discussion of four interrelated goals of MAD. Current gaps in MAD are identified from the data and information, processing and systems perspectives. The selected literature review is structured around specific organizations known to be active in maritime anomaly detection, various MAD systems, and other relevant research activities. Finally, Part 4 evaluates concrete anomaly detection techniques presented in Part 3. Based on this assessment, recommendations about future directions are made.

To prevent confusion, the use of the terms “data” and “information” must be clarified. Some people consider them synonymous, but most of the time, information is considered to be distilled data, and therefore is judged as more valuable. For the purposes of this paper, the term “data” is used to refer to both concepts, unless explicitly stated otherwise.

2. The universe of maritime anomaly detection

There is no consensus on what maritime anomaly detection is. In fact, the community is very far from agreement on several of its aspects [DARPA, 2005-G]. When it comes to terminology, nomenclature, processes, needs or issues, the various players in the community have their own vision of the problem. A very broad definition that seems to fit almost all visions could be the following: finding unusual behaviour in the maritime domain and evaluating its threat potential.

Although the actors in this domain use very different approaches to handle the problem, they all share the same objective: to exploit maritime domain information resources to improve people security. The following sub-sections summarize the essence of maritime anomaly detection. First, the major steps are presented as a process used to concretize MAD. After that, the goals to be achieved through the application of this process and, finally, a list of technological gaps that prevent the realization of these goals are provided.

In the following three sub-sections, the different views of the problems of MAD are synthesized. Each actor addresses, in its own way, a piece of the puzzle. These pieces are summarized and merged to provide the reader with an overall picture.

2.1. Steps in maritime anomaly detection

The detection of maritime domain anomalies is obviously not a monolithic process; it can be easily broken down in several ways. We have chosen to divide the process into seven parts. That decision is not arbitrary. First, the process appears to be generic, i.e., not tied to the maritime domain. In fact, the detection of air or land anomalies could be broken down in the same way. Second, the process as presented by different authors maps easily to the breakdown presented here. Finally, these steps go beyond simple detection. By presenting threat assessment, for example, which is always included in the process anyway, we can go beyond the basic detection process by including the question "So what?" that is so important to analysts.

2.1.1. Data and information acquisition

Data is the raw material of anomaly detection; it is therefore logical to make data acquisition (or collection) the first step of the process. Anomalies are derived from data, so the success of the following steps depends largely on data collection. If the data is absent, insufficient or of poor quality, all the following steps will be compromised. If there is no data, there will be no anomalies.

In recent years, systems have been developed to collect as much information as possible. They were not (or were rarely) meant to be used in anomaly detection, but rather to raise situational awareness. Sensors provide data on a variety of aspects of the maritime domain, such as the positions of ships, weather and tides, and this information is so abundant that it is difficult to handle it properly. The problem of being drowned in data is referred to as "data overload" [Kessler, 2009] and it is the reason why anomaly detection exists: to sift through large quantities of data and highlight elements worthy of interest.

Systems are always designed to increase awareness of the real world by using their sensors to construct digital representations of it [Boner, 2009; Boraz, 2009]. The evolution of computers in recent decades contributed greatly to data sharing and storage. Nevertheless, data acquisition is not a simple task; to reach goals, various data sources, public and classified, must be used [Kessler, 2009], and they may be of different pedigrees, be geographically isolated or reside in "stovepiped" systems that prevent cross-organization sharing [Moore 2005].

Part of anomaly detection can be done at this stage. It consists of identifying inconsistencies or extreme values in the data. For example, a ship with an Automatic Identification System (AIS) transmitter that spontaneously stops in the coverage area of the receptors may be interpreted as an anomaly. A transmitter may have failed, the ship may have sunk, or illegal activities may be being camouflaged [Baldacci, 2008-A].

2.1.2. Data and information fusion

One can view the data acquisition process as a tentacular process that reaches for many different sources. Often, the retrieved data concern the same subject. For example, if one wants to know the temperature in New York City, one can ask people around the city to take a reading from their thermometer. One may, and probably will, receive different values. In the fusion process, all these values are merged to provide the most accurate value as possible.

Data and information fusion is a very active research topic. It started with the sensor community and it now covers the higher-level aspect of data and information. Fusion is a key enabler in achieving high-quality situational awareness; that is why it is the second step in the anomaly detection process. A great deal could be said about fusion, but that is beyond the scope of this document. Fusion is one step in the process, and the task of exploring it in greater detail is left up to the reader. However, its relationship to maritime domain anomaly detection is briefly described.

Various sources provide data of different pedigrees, format and precision [Kessler, 2009]. Moreover, those sources are not necessarily reliable and may provide incomplete or uncertain data [Walden, 2006]. The challenge of fusion as a process is to associate, correlate and combine the data from multiple sources, taking into account all these factors, and report fused information that is as accurate as possible. As noted in [Kessler, 2009], sometimes, for reasons of classification, data fusion of classified and unclassified data cannot be performed. This is necessary to satisfy security restrictions defined by the classification level of the fusion process recipient.

The position of ships is a classic example of fusion in the maritime domain. There are several ways to locate a ship on the ocean, but accuracy and reliability vary greatly depending on the data provided by sensors. Radar contacts come with an error ellipse and are limited to line of sight. Positions reported by reconnaissance aircraft and unmanned aerial vehicles (UAVs) can be imprecise. AIS emissions can be forged or stopped. The fusion process must take all these constraints into account and provide the user with a geographic position as close to reality as possible.

Sometimes, a large discrepancy between two or more sources in the fusion process can be considered as an anomaly. If an AIS notification indicates a position but the triangulation of the

transmitter is different, that could indicate an attempt at deception, e.g., a forged AIS transmission [Smith et al., 2009].

2.1.3. Situational awareness

Data of sufficient quantity and quality can provide a fair representation of the current state of reality. If the data cover all aspects of a situation of interest in a timely manner, one can then say that complete and continuous situational awareness has been achieved. This "picture" of reality is the basis for effective decision making. In addition, it facilitates understanding from a higher level of abstraction [Boraz 2009] and understanding of the interactions between different entities [Kessler 2009]. However, complete situational awareness would be akin to omniscience and achieving it would be a utopia. When this perception is deficient, it can lead to serious interpretation errors.

At this level, most information (and knowledge) is derived and not directly observed or reported [Kessler 2009]. The previous step (fusion) has reduced the volume of data and improved data quality. One might be inclined to believe that having data on different aspects of reality should lead to greater enlightenment concerning the situation, but that is unfortunately not always the case. The information can still be too vast to be fully assimilated by an operator; the presentation step will handle this problem.

There are currently systems that are designed to support the operator in reaching a certain level of situational awareness. These systems typically track vessels on a world map displayed on a screen. Simply following the movement of ships on a map can help the operator understand some of the behaviour in the maritime domain, but that alone is not sufficient to explain the behaviour, which is the actual goal. That being said, with technological advances, it is possible to know the position of vessels, their speed, and cargo owners, just to name a few aspects. There seems to be unanimous agreement on the need to add inputs to these systems in order to increase the operator's awareness. In fact, it is the solution proposed in almost all the documents; the operators want to fill in the information gaps to improve their situational awareness [Moore, 2005].

For example, a loaded merchant ship loitering or anchored in an unofficial anchorage area may arouse suspicion. However, the boat may be ahead of schedule and perhaps the captain does not want to pay harbour fees. Another possibility would be that it is waiting for a more favourable market for its cargo before unloading [Tarsus, 2009]. An operator with a good level of awareness would be able to judge this situation and, if necessary, raise an anomaly.

2.1.4. Anomaly detection

When the understanding of the maritime situation reaches a certain level, it may be possible to detect certain patterns. These patterns may not always exist, but studies have shown that, most of the time, they do [Seibert, 2009; Moore, 2005]. By definition, a pattern is composed of recurring events that repeat in a predictable manner. Here, what is predictable is considered normal and the rest will be labelled as anomalies. That is one definition of normality; there are many others that do not totally agree with each other. In fact, there is no consensus on the definition of the term "anomaly," and that causes some confusion about what is normal or abnormal [Baldacci, 2008-A]. In the literature, several terms are used: "outliers," "anomalies," "unusual," etc. One

way to sum up all of these definitions could be as follows: a perspective in which an observation does not seems to belong to any group. That definition is, of course, debatable.

In this step, a model to discriminate normality is constructed using data collected in the past (i.e., historical data). It is a pattern recognition step, and it can be goal-driven or data-driven. Pattern can be, for example, a sequence of events, a statistical distribution, or a cluster of elements. Subsequent data will be compared to the model to classify them as normal or abnormal. If the pattern is not subject to change over time, the model is said to be static. The model can also be dynamic if the patterns evolve over time.

There are many methods for discovering patterns. The choice of a particular algorithm depends on a number of factors. The data format, the performance requirements and the nature of anomalies are some considerations that influence the choice of a particular method. Describing the techniques used to detect anomalies is a colossal task that is beyond the scope of this document. However, the main families of solutions, taken from [Hodge et al., 2004], are presented here in order to lay the groundwork for later sub-sections of this document where the practical applications of anomaly detection are presented.

2.1.4.1. Statistical methods

There are two types of statistical techniques: parametric and nonparametric. With parametric techniques, if the data correspond to a particular statistical model, anomalies can be detected rapidly and without supervision. With nonparametric techniques, no assumption is made about the underlying distribution of the data. Although more resources are required to develop them, these methods are effective for automated anomaly detection.

Statistical techniques are simple to implement, but their capability is limited to specific problems. Ship speed is a good example of a variable with which these techniques are effective because the anomalies are extreme values. In cases where anomalies are uniformly dispersed in the sample, these techniques are ineffective. Moreover, since it is difficult to define a threshold for separating abnormal values in a normal distribution, statistical techniques are likely to have a high level of false positives.

2.1.4.2. Neural networks

There is a great deal of literature on neural networks, and they come in many varieties. Overall, we can say that they generalize well to unseen patterns and are capable of learning complex class boundaries. After training, the neural network acts as a classifier. However, all data must be traversed several times before the network converges to an appropriate data model. Training and testing are also required in order to fine-tune the network and determine threshold settings before neural networks are ready to be used for classifying new data.

One of the biggest criticisms of neural networks is that the process is very obscure; they are often referred to as black boxes. The processing between the input and output neurons is not intelligible and cannot provide the operator with explanation or justification. In

addition, they are subject to overfitting: if the learning phase is too strict or targeted, classification performance on data near a class boundary may drop.

2.1.4.3. Machine Learning

Machine learning is not a technique but rather a field of research. It is a candidate of choice for anomaly detection. The main focus of the discipline is to automatically learn complex structures and make decisions based on the data. This focus is similar to that of an anomaly detection system.

Several documents mention that the use of machine learning is desirable, but little detail is provided. This reflects the diversity of problems in anomaly detection. Indeed, there is no single technique that could meet all the requirements; restrictions typical of each situation can be very different. Here is a short list of popular techniques belonging to this field: decision trees, genetic programming, support vector machines, Bayesian networks, clustering, etc.

2.1.5. Anomalies put in context

Context plays an important role in anomaly detection. Because patterns used to detect anomalies cannot take into account all environmental factors, it is necessary to put each anomaly, once detected, in context. This "side" information can be used to justify the behaviour of an entity. This is one more reason why good situational awareness is needed to explain an event. Relevant contextual data qualify the anomaly detections [Seibert, 2009]. A deviation from a routine behaviour can be a simple response to environmental conditions [Kessler, 2009], such as a hurricane or an iceberg. In some cases where no anomaly is raised by pattern matching, context can raise suspicion. Benign behaviour that seems perfectly normal may be an anomaly when it is put in context, especially in cases of deception or covert behaviour [Sisk et al., 2009].

For example, a system that observes a ship and looks for anomalies in its velocity could conclude that its speed is consistent with average speeds of other ships travelling the same route, and that everything seems normal. Although the speed may be well beyond the normal speed for this type of ship, since the system does not take that information into account in its anomaly detection routine, everything seems normal. Another example: the system may raise an anomaly for a ship stopped in a sea lane. Since the system does not know that it is a research vessel and that the behaviour is normal, a false alarm is raised and will have to be taken care of by the operator.

One can see here the importance of situational awareness for putting anomalies and entities in context. One should leverage the power of computer systems to detect as many anomalies as possible but, in the end, a human must make the decision. For an operator to pick up the task where the computer system has left off, he or she must possess contextual information in order to make the right final decision.

2.1.6. Threat assessment

Abnormality is not synonymous with threat. Sometimes in the literature one cannot always distinguish the two. There is a simple explanation for this: the two concepts are fused. For some, an anomaly is only a threat in context [Seibert, 2009]. Some people do not care about anomalies that present no threat, and often all that matters is the threat to blue forces [Sisk et al., 2009]. Anomaly detection is just a tool to detect threats.

The definition of a threat is very broad and encompasses a multitude of aspects. Terrorism is often considered a threat, but threat assessment should include all activities that could harm a nation. Here are some activities that are generally considered to be threats:

- Terrorism
- Illicit traffic (weapon, drug, WMD)
- Spying
- Piracy
- Illegal fishing
- Military manoeuvres
- Territorial violations
- Pollution

The ability to provide information on the threat level is a value-added for systems. The ability to classify the anomalies under the labels "benign," "not explained" or "threat" can greatly enlighten the operator [Seibert, 2009]. It is also desirable to quantify threats so that they can be handled in order of priority.

In several concepts and/or systems, anomaly detection is a means for arriving at an automated detection of unanticipated threats [Sisk et al., 2009; Boner, 2009; Moore, 2005; DARPA, 2005-D]. It is worth mentioning that there is also a demand for preventive detection of potential threats. Without presenting any anomaly, a vessel may have the capacity to cause much damage or may have a valuable cargo coveted by pirates. These vessels require special attention and can be identified as "Vessels of Interest" [Barrett, 2009].

2.1.7. Dissemination and presentation

The final stage of the detection of anomalies is to share analysis results. The outputs of a system can be multiple and of different formats depending on the recipient. Reporting to humans is a science. Presenting the big picture to increase situational awareness is related to visual analytics, and that involves more than just blips on a monitor [Boraz, 2009].

Human/computer interaction plays a critical role in understanding and evaluating anomalies [Griffin et al., 2009].

The dissemination of information to various non-human entities is also of crucial importance. Several systems can reuse the output of different steps of anomaly detection, or be informed of the conclusions of the analysis. There is no need for a system to be monolithic and, in fact, collaborative distributed systems are envisaged.

As previously mentioned, the operator's work begins where the system's works ends. Therefore the design should provide users with the best possible situational awareness to enable them to make the right decisions. Clear and concise presentation of explanations of the alerts generated by the anomaly detection process is an important challenge [Moore 2005].

2.2. Anomaly detection goals

To understand why so much effort is devoted to research and development in the area of anomaly detection, it is necessary to know the intended purpose of AD technology. As mentioned at the beginning of this document, the main purpose is the security of the population. However, systems already exist and the research in this field focuses more on improving existing technologies. This section reviews what is expected from current research efforts.

2.2.1. Manpower optimization

The number of vessels circulating on the surface of the globe keeps increasing. Meanwhile, the staff required to oversee the maritime domain is being reduced [DARPA, 2005-A]. Moreover, the increasing terrorist threat of the last decade has raised performance expectations for security systems. This creates an obvious contradiction: one must do more with less. To solve this issue, the productivity of our workforce must be maximized.

To achieve that objective, it is preferable to use the complementary strengths of humans and computers. Machines are capable of processing huge quantities of data quickly; humans are not. Human reasoning capacity is far superior to that of the machine. Therefore, the proposed approach is to let the system do the simple, routine reasoning and to steer the attention of operators to more complex problems [Walden, 2006].

2.2.2. Support decision processes

Human are always in the loop. Systems are made to improve the performance of operators, not to fully replace them. The amount of data that enters a system is typically astronomical, and a single person cannot manage and interpret it quickly [Walden, 2006]. For this reason, when an operator needs to make a decision, he/she must possess all the

relevant information to orient his/her thinking. Systems must allow users to reach a certain level of situational awareness through proper presentation [Moore, 2005]. For that, all steps of anomaly detection must be executed.

These steps are summarized from the perspective of a d ecision maker's need, highlighting the usefulness for his/her task. First, the relevant data are collected from a multitude of sources to be merged and cleaned. With this data it is possible to get an overview of the situation where the anomalies are detected and put into context. Finally, the decision maker will be provided with the situation, the anomalies and their justification. Decision makers are accountable for their actions. That is why this type of support will help them make the right and justifiable decisions based on facts and not rely on intuition or luck.

2.2.3. Predictions and early notification

It is desirable to be able to predict negative events. This kind of capability helps authorities prevent such events from occurring or, if that is impossible, to at least prepare for them. Also, to be useful, these predictions must be made as early as possible to give time for the authorities to distribute information about the situation. For example, a ship travelling at full speed in the wrong direction in a waterway could be a potential collision risk. An anomaly should be raised as soon as the ship going the wrong way is detected, not when an actual collision occurs.

In fact, for every threat that eventually materializes, there is precursor behaviour that deviates from normality, identifiable after the event. The goal is to provide early detection of these signals before the event. For example, a presentation to the industry of the Defense Advanced Research Projects Agency (DARPA) project Predictive Analysis for Naval Deployment Activities (PANDA) [Moore 2005] uses an attack on an American vessel by a merchant ship as an example. The merchant ship is hijacked and then a rendezvous at sea is carried out. Following this first anomaly, two major deviations from the normal route are made before the attack on the U.S. Navy's Kitty Hawk CVG takes place. In this example, there were three opportunities to detect an anomaly, and an early report could have prevented the negative actions. These kinds of situations are cited as examples to justify the need for anomaly detection systems.

2.2.4. Maintain a complete and continuous surface picture

The final point in the goals sought in the maritime domain is the ability to monitor everything at all times. It is difficult for a human being to maintain a constant level of concentration and impossible for him/her to pay attention to everything. However, that kind of task is the strength of computer systems, as they provide constant performance, their reasoning is deterministic and they have an exact memory with a large capacity.

Thanks to this observation, it is easy to orient the development of anomaly detection systems. The goal is to almost completely automate surveillance in order to achieve a constant overall situational awareness. As much data as possible is already collected to cover the entire maritime domain at any time, thereby causing the problem of data overload. Systems will process this data to maintain a constant and complete surface picture and will notify the operator only in the event of anomalies.

2.3. Anomaly detection gaps

To achieve the identified objectives and follow the steps in the detection of anomalies, several technological gaps must be filled. Up to this point, a rather idealized solution has been presented to explain the actual goals of anomaly detection, and the problems that need to be solved to reach these goals have been set aside. The common technology gaps identified in the literature are now presented and placed in the context of the maritime domain.

2.3.1. Data and information

Access to data and information is critical; data is the basis of the overall process of anomaly detection. Therefore, relevance, abundance and diversity are increasingly sought. However, several shortcomings have been identified so far. The information gaps that prevent us from achieving adequate situational awareness and detecting anomalies effectively are discussed next.

2.3.1.1. Geo-spatial positioning of vessels

Knowing the position and kinematics of vessels is essential for all planning in the maritime domain. This capability, called "tracking," is relatively well developed. However, as noted previously, one can always use more information. In fact, tracking is based on the fusion of multiple sensors such as imaging, radars and AIS transmitters. In their current state, such technologies are not able to provide comprehensive global coverage of the maritime domain.

The coverage of AIS receivers is often limited to certain coastal areas, and only vessels over 300 tons must be equipped with a transmitter [DARPA, 2005-A]. Moreover, it is rather easy to stop the transmitter [Baldacci, 2008-B] or to falsify a report [Ristic et al., 2008]. The radar is also limited in range and can miss small vessels in a cluttered environment [Smith et al., 2009]. Imaging is far from being able to cover the entire world globe and cannot accurately identify the vessels.

2.3.1.2. Cargo, ports and owners

The control of cargo transition is essential to the security of a country. Cargo ships can constitute a threat with weapons of mass destruction, other weapons, drugs or other illegal items on board. Having access to the ship's manifest is a start, but manifests can be falsified. Therefore it is important to know who owns the vessels, so that links between maritime transport and criminal organizations can be identified. In addition, certain countries are known for their lax control of goods, and criminal organizations and terrorists can exploit these vulnerabilities.

This information is not always directly related to the detection of anomalies. Certainly, if a vessel carrying Liquefied Natural Gas (LNG) has as its destination a port that cannot receive gas, an anomaly should be raised [Tidepedia, 2009]. But most of the time, this type of information will be used for threat assessment. Indeed, a ship carrying weapons may behave quite normally in order to avoid attracting the attention of authorities. However, the manifest, the owner and the ports visited may raise suspicions. Currently, only 50 percent of global shipment is covered. Many records lack key elements, and sources vary in reliability [Boner, 2009].

2.3.1.3. Passengers and crews

As well as the cargo, passengers and crews may represent a threat that cannot be detected just by watching movements of ships. An analogy can be drawn here with air transport passengers. Each passenger has a different background, and information about it can be used to identify threats to the interests of a country. Even if the threat is not related to the maritime domain, there is an opportunity to identify and handle it.

Although it is desirable to have a complete profile for each person on a ship, it is not realistic for now. However, efforts should be made to fill this major gap that could make a significant difference in threat assessment. Basic information, such as name and nationality, could be used to retrieve and build a more complete profile of a person using external intelligence databases. Here is a short list of what the system may be looking for regarding passengers and crews: terrorists, criminals, refugees, diseases, etc.

2.3.1.4. Economical and sociopolitical factors

Some anomalies may be justified by global economic conditions or by cultural or political factors. The fluctuating price of oil influences the cost of transportation of commodities, which in certain circumstances can change the type of cargo or the volume of traffic on a sea lane. There is a need to exploit the global maritime trade data to formulate a model of what goes in and what goes out of major economic centres [Boner, 2009]. In addition, there are invisible barriers such as zones of exclusion, bans, embargos and fishing that must be taken into consideration when the position of vessels is analyzed. For the moment this kind of information is lacking.

For example, if the price of oil goes down enough to make the price of Brazilian bananas attractive to the French market, a sea lane could arise between Brazil and France. However, the volume of bananas between Spain and France will probably decrease. As another example, the presence of an American merchant ship in a Cuban port is clearly a violation of the embargo imposed by the U.S. on Cuba.

2.3.1.5. Environmental factors

Of all the elements necessary to understand a situation, environmental factors are probably the easiest to justify. That may be because most people watch the weather to plan their day. The same is true for navigation: weather, tides, icebergs and other natural phenomena greatly influence shipping [Seibert, 2009]. It is not uncommon for a ship to deviate from its route to avoid crossing the path of a hurricane [Tarsus, 2009].

Large quantities of environmental data are currently collected from different organizations and agencies. The problem is that this information is often in "stovepiped" systems and is virtually impossible to use in newer systems. Also, there is always a need to gather new types of information to improve awareness. There are currently demonstration systems that collect some weather data from collateral databases.

2.3.2. Processing

Owning a lot of information is desirable, but one must be able to exploit it. As mentioned in previous sections, an abundance of data is available and the processing capabilities are insufficient to synthesize and present it in a meaningful way to the user. The following subsections discuss areas where the processing capacity has been regarded as insufficient.

2.3.2.1. Knowledge base and ontologies

In order to partially replace a domain expert with an automatic system, one must emulate his/her capabilities. The problem lies in being able to capture the problem-solving knowledge gained through experience over time by the expert and then encode that knowledge in a system. Since many entities work and refer to the same domain, this encoding of knowledge must be done in a common language.

Knowledge bases and ontologies (e.g., a Maritime Information Exchange Model) should be constructed to develop systems that can reason, communicate and interact [Griffin et al., 2009]. The more complete the knowledge support, the greater the ability to understand the maritime domain and to provide support for the process of anomaly detection.

2.3.2.2. Natural language processing

The use of structured data is very common today. However, a huge portion of the relevant data is in unstructured formats. One need only think about the Internet, paper reports or conversations between two individuals. These information supports are meant to be used for communication

between humans. Machines have their own way of communicating and have a very limited capacity for understanding natural languages. Moreover, humans communicate a lot, in many different languages or dialects, and we use special constructs such as sarcasm. From an information theoretic point of view, human communication is a fantastic contraption.

Humans do not possess the necessary resources to process all these communications. Moreover, the fact that we do not know the content of these communications prevents us from selecting the relevant ones for processing. Therefore, it is necessary to advance the processing of natural languages in order to automate this process, even if only partially [Griffin et al., 2009].

2.3.2.3. Network analysis

The world in which we live is becoming more complex and interconnected than ever. The communication capabilities provided by the proliferation of electronic devices allow us to interact quickly with a multitude of individuals who are separated geographically. That capability has been constantly increasing in recent years; the telephone, the media and now the Internet have all contributed greatly to it.

The power to understand, explain and influence the structure of relations between different entities is a major asset for detecting anomalies. One need only think about terrorist networks or criminal affiliations to grasp the importance of this concept. In practice, in this case, one must not model the normal behaviour and find entities that do not conform to it. One must model deviant behaviours. At the moment, the tools or the knowledge of these threatening organizations is insufficient to develop this capability.

2.3.2.4. Hypothesis creation

Anticipation and creation of different scenarios are cognitive processes that facilitate the understanding and management of a situation. An entity that is a threat to a state has, by definition, a course of action that might affect the state's integrity. Creating hypotheses can help us anticipate anomalies, assess the threat and provide a list of possible actions to mitigate or prevent a harmful situation.

Being able to provide a list of probable sequential events following a situation is a major asset for the decision maker. This ability is similar to asking the opinion of an expert on the possible outcomes. Hypothesis creation can be useful at several stages of anomaly detection and is thus a desirable asset. However, hypothesis creation is an extremely complex problem because it is difficult for both humans and machines to sift through all the possibilities.

2.3.2.5. Motion-based pattern learning

Like a system operator who monitors maritime traffic, a computer system should be able to learn to discern the usual traffic. According to studies, between 75% and 85% of all maritime traffic has a structure discernible by a human being [Seibert, 2009]. It is therefore possible to automate the modelling of major sea lanes for the purpose of anomaly detection.

Prototype systems that have this kind of capability already exist. However, there is a strong demand to push the limits of current technologies [DARPA, 2005-F]. The goal is to quickly sift through a large number of ships and discover new structures, even though very little data is available. For example, the final objective of the PANDA project is to be able to detect a normal trajectory with less than six samples in 90% of cases with an error rate of less than 1% [DARPA, 2005-B].

2.3.3. Anomaly detection systems

Several anomaly detection systems already exist and are in operation. However, some are just simple prototypes. The lessons learned from these current systems allow us to consider the desirable features for future systems and identify needs that are not covered yet. The following are a few areas in which development efforts are currently required.

2.3.3.1. Policy restrictions

Restricted access to certain data and information is a major problem in the domain of anomaly detection. There are a multitude of systems and agencies that possess relevant data and information, but gathering some of that knowledge is a serious challenge. Moreover, different security levels are a limitation in data fusion [Kessler, 2009].

Access restrictions are of course justifiable, but it is necessary to develop a system that could benefit from all these sources of data. The ideal system should have access to all sources of data and know their level of confidence. Such a system would make the presentation to the user after taking into account its level of security and the need to know.

2.3.3.2. Data handling

With the increasing number of data sources also comes a need to adequately store the information. Data archiving is useful when searching for an event or when one wants to replay a whole situation. The ability to perform data mining and domodelling tasks (for automated reasoning) is also desirable.

All these technologies already exist separately. However, the challenges lie more in their cohabitation. The archiving process moves the data into permanent slower and cheaper storage that makes the data harder to access. For data mining and modelling, rapid access to historical data is needed, and this is where conflicts arise. In order to maximize data exploitation, one must find a way to keep information without compromising speed of access.

2.3.3.3. Support requests from many users or partners

As previously mentioned, sharing data and information among organizations is desirable. It prevents duplication of collection and storage processes. Along the same logic, sharing processing resources could also be beneficial. Legacy systems are often criticized as being closed and monolithic. Because resources must be fully exploited, future systems must allow multiple

users to send requests to various services. Here, the term “user” refers not only to entities within the same organization, but also to external partners.

Services-oriented architectures (SOAs) are a good way to share the resources of a system. These solutions seem to be very popular, and their use is strongly suggested in most of the recent literature. The gap here is the need to develop new open systems and retrofit stovepiped systems so that they can share their resources.

3. Selected literature review

A review of some relevant work is provided in this part. Each section describes an article, a presentation or a lecture about a subject related to anomaly detection and focuses on the inner mechanics of the methods and the assessment of results.

3.1. Research at the NATO Undersea Research Center

The North Atlantic Treaty Organization (NATO) Undersea Research Center (NURC), based in La Spezia, Italy, is a research organization under NATO's Allied Command Transformation. NURC conducts research and develops products to support NATO maritime operations and to support the continuous improvement of NATO military capabilities. One of the research areas at NURC is maritime situational awareness and the goals are to develop new techniques and technologies for monitoring the world's shipping channels. Some of the results of this research are presented next.

3.1.1. AIS coverage estimation

The Center of Marine Sciences (CCMAR) made a request to NURC for access to short-term (6 to 12 hours) and medium-term (4 days) coverage maps of the AIS contacts. Although they do not contribute to anomaly detection, recently updated AIS coverage maps are of great help for

- multi-sensor allocation, tasking and data fusion;
- asset scheduling and resource management;
- AIS-based anomaly detection algorithms;
- completing the maritime picture from an operator's point of view.

Traffic-based AIS coverage [Baldacci, Fabiani, 2008] depends on two factors: a) the actual traffic; b) AIS status transitions, i.e., the arrival of new AIS contacts ("birth process") and the loss of AIS contacts ("death process"). In this context, a contact is declared to be lost when it is not updated by a new message within a maximum allowed interval. The first step in the building process is to divide the area of interest, in this case the Black and Mediterranean seas, into 0.1 by 0.1 degree cells. Next, the AIS coverage index is calculated by taking into account all the contacts falling into the cell. In order to have meaningful maps, the coverage index is averaged over periods of at least few hours (typically 6 to 12) and then a closing morphological operator is applied in order to obtain maps that are more continuous in space. The final product is put on the NURC Open Geospatial Consortium (OGC) compliant map server and can also be viewed with Google Earth.

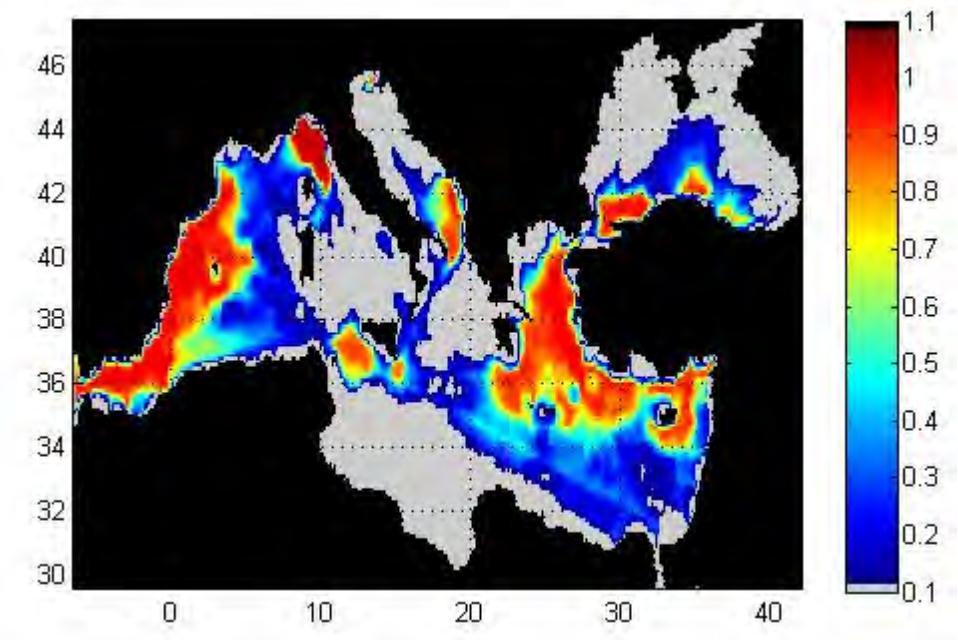

Figure 1 - Medium-term AIS coverage estimated over March 2007 [Baldacci, Fabiani, 2008].

3.1.2. Vectorial maritime traffic characterization

Many methods have been tried for modelling maritime traffic. A vectorial traffic characterization is proposed in [Baldacci, Carthel, 2009]. The reasons for this approach are based on conclusions from previous work indicating that

- bottom-up approaches are usually preferred for traffic characterization;
- among the bottom-up approaches, grid-based ones have proven ineffective so far;
- Gaussian mixture modelling is very difficult due to the possible combination of many components;
- users are interested in obtaining an analytical description of the traffic.

Based on these conclusions, the authors consider a simplified version of the problem: traffic characterization of main sea lanes only.

To extract vectors corresponding to sea lanes, AIS contacts and many image-processing techniques are used. Because of the complexity of the procedure, the complete mechanism will not be explained here. The AIS contacts are plotted on a map and, after a filtering pass, an edge-detection algorithm is applied to extract segments of sea lanes (see Figure 2). A clustering pass is then made on these segments, and only one segment describing each cluster is kept.

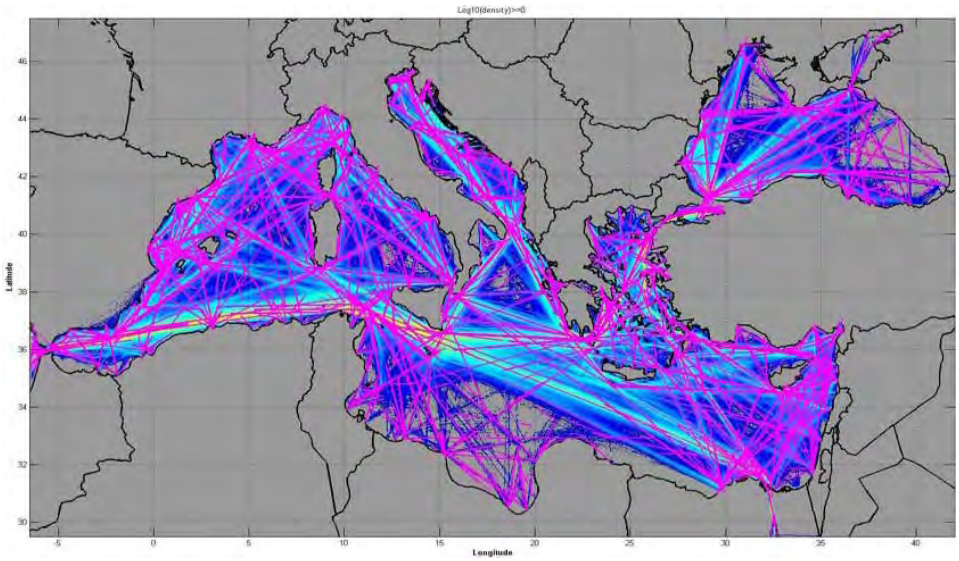

Figure 2 - Extracted sea lane segments [Baldacci, Carthel, 2009]

Each resulting segment will form a sea lane region characterized by a Gaussian distribution, and ships will be associated with it based on their distance from the corresponding segment. Vessels in these sea lane regions are then submitted to an anomaly detection test using the normal distribution for course over ground and speed over ground.

3.1.3. AIS transponder anomaly detection

When monitoring the AIS traffic, the main events of interest are the arrival of new AIS contacts (birth process) and the loss of AIS contacts (death process). Among the two types of events, AIS deaths deserve more attention, as ships can intentionally turn their AIS transponder off in order to hide their—potentially illegal—activity. [Baldacci et al., 2009] present a technique to sort different types of AIS death to trigger the attention of an operator on suspect cases.

A model of the AIS birth and death processes is built based on contacts provided by the NURC Maritime Mobile Service Identity (MMSI) server. This server provides contact data every five minutes and eliminates all duplicated entries. For each vessel, the probability of a contact is calculated using the following:

- The probability of receptor failure (0 or 1).
- The AIS coverage map index (as in [Baldacci, Fabiani, 2008]).
- The behaviour of the transmitter (calculated using a time series).

This probability is then used in one of the two proposed statistical hypothesis tests, Neyman Pearson or maximum *a posteriori* (MAP), to detect an anomaly.

The logic of this algorithm can be summarized as follows: On the one hand, whenever the probability is high (i.e., in areas with good AIS coverage and for reliable transponders), if a contact is lost, the death is more likely to be due to an emission termination and therefore only a

few observations are necessary before an anomaly is declared. On the other hand, when it is low (either because the ship is moving out of the AIS coverage area and/or because the AIS transponder is not reliable), it is necessary to have more observations, i.e., to wait for more scans before declaring an anomaly.

3.1.4. AIS anomaly simulator

In order to develop and train the intelligent agents and to assess their performance, the ground truth is needed. For this reason, NURC has developed an AIS anomaly simulator in the Matlab programming language [Baldacci, 2008-B]. With this tool, it is possible to simulate AIS status, kinematic and positional anomalies.

This tool uses the AIS emission model described in [Baldacci et al., 2009-A] to construct time series of AIS emissions from tracks. The tracks can be random or deterministic, and in the later case the input is made via an XML file using waypoints. In order to simulate the perturbations in speed over ground and course over ground, which typically affect real navigation, random perturbations are added to the computed values. Additionally, AIS anomalous tracks can be injected into the Maritime Surveillance Data Simulator (MSDS) developed at NURC and used to evaluate the performance of multisensory data fusion and multi-sensor anomaly detection algorithms.

3.1.5. Analysis of AIS intermittency using an HMM

[Guerriero et al.] describe an AIS on/off detector that relies on a higher-fidelity model of the AIS transmission channel than that adopted during the *Maritime Surveillance 09* (MS09) sea trials. In particular, it introduces the notion of channel memory through the mathematical formalism of Hidden Markov Models (HMMs). This allows a disambiguation of dropouts due to channel effects and dropouts due to suspicious vessel behaviours.

One of the main objectives of the MS09 sea trials was to test algorithms for AIS anomaly detection. The simplest form of AIS anomaly is missing transponder measurements in regions of good AIS coverage. Emission termination may imply that a vessel has turned off its AIS transponder. During these trials, the AIS coverage intermittency had a significant impact on the AIS on/off detector tested onboard, which did not properly account for the characteristics of the AIS transmission channel.

It has been observed from the trials that the AIS channel might be bursty or intermittent. To model this behaviour, a simple two-state HMM is used where one can represent the states as a valve that can be open or blocked at any time. The same method is used to model the vessel's AIS emissions, i.e., on or off. These two models are then fused in a four-macro-state HMM that can be used to detect an anomalous use of the AIS transponder. However, the channel is found to have high memory (i.e., burst, prolonged dropouts), leading to a nontrivial false-alarm rate for reasonable detection performance.

3.2. Research at BAE Systems

BAE Systems is a British defence, security and aerospace company headquartered in Farnborough, Hampshire, England. BAE is one of the world's largest defence contractors. It conducts research in maritime situational awareness and is a partner in the DARPA PANDA project. This subsection provides a review of some of BAE's publications.

3.2.1. Maritime situation monitoring and awareness using learning mechanisms

[Rhodes et al., 2005] address maritime situational awareness by using algorithms to learn behavioural patterns at a variety of conceptual, spatial, and temporal levels. Continuous learning enables the models to adapt well to evolving situations while maintaining high levels of performance. The learning combines two components: an unsupervised clustering algorithm, and a supervised mapping and labelling algorithm.

At the core of the system lies a significantly modified version of the fuzzy ARTMAP neural network classifier. The learning system must initially be presented with a series of observations that are labelled as normal/acceptable by a subject matter expert. After this first phase, the system can classify events on its own. If an event is not covered by the classifier, an alert is raised and the operator can classify it manually. Alerts ignored for too long are assumed to be benign events and are labelled as normal. Multiple demands can be made on human participants via alerts so the system sorts them in order of distance from existing clusters.

While no performance comparison tests are provided, it is claimed that despite limited initial bootstrapping data it is still possible to achieve a well-performing model with minimal operator effort. Two academic examples using real data are provided.

3.2.2. Adaptive mixture-based NN for automated behaviour monitoring

[Garagic et al., 2009] proposes replacing the deterministic fuzzy ARTMAP hyperrectangular mapping discussed in [Rhodes et al., 2005]. This previous method assigned uniform probability inside the hyperrectangle and did not consider the distribution inside the category. To improve their method, the authors propose using an adaptive mixture of multidimensional probability density component.

The adaptive mixture-based neural network classifier algorithm is composed of three main steps. The first step determines the highest *a posteriori* probability to assign to a category. In the second step, the Mahalanobis distance (distance based on correlations) between the observation and the chosen category is calculated. If the distance is too high, an anomaly is declared. In the other case, a new category can be created or the probability density function is updated using the expected maximization and the Kullback-Leibler information metric. The last step tries to merge the closest existing categories.

Despite the fact that no real performance test is provided, it is claimed that this method outperforms previous ones. It is supposed to be robust to noise and operator error. The speed is

fast enough for real-time applications and it maintains a high level of fidelity in separating normal and abnormal behaviours.

3.2.3. Associative learning of vessel motion patterns

The work presented in [Bomberger et al., 2006] describes a learning-based approach to providing predictions of future vessel location given the actual position and velocity. The system associates different geographical grid locations through Hebbian learning corresponding to the position of a vessel at constant time intervals. The time intervals for prediction can be selected to suit the operational needs of the users.

The concrete implementation places a uniform square grid over the area of interest to discretize the vessel location. It also defines a discretization of the vessel velocity that enables learning to be contextually specific to the behaviour of the vessel. Thus, for each vessel report, it is able to place the vessel in a grid location and give it a velocity state. Weights are attributed to pairs of grid locations/speed and change via Hebbian learning. They will decay over time but will be reinforced by the passage of a vessel. To detect an anomalous vessel position, the system uses the previous grid position, and if the weight is not strong enough, an alert is raised.

Performance of the system was tested with AIS data from the Miami harbour, and the correct prediction rate never rose above 70 percent. Also, prediction accuracy is greater for arriving vessels, and the system is completely ineffective in the open sea. The learning process seems to require a whole month's worth of data in order to achieve maximum performance and coverage.

3.2.4. Multiple scales probabilistic associative learning for MDA

Improvements were proposed in previous work [Bomberger et al., 2006]. These improvements were implemented and results are presented in [Rhodes et al., 2007]. One of the problems was that the predictions were inaccurate in the open sea, and it was suggested that that could be the result of an improper grid size. Also, the predicted areas were selected from an arbitrary threshold on the weight given to destination location in the Hebbian learning.

To solve the first issue, tests were made using different grid sizes. The results showed that different grid sizes should be used depending on the zone around the harbour. Four zones were created: Miami River, Miami harbour, the controlled approach area east of the harbour, and the open water (Fig. 3). The optimal grid size was computed for each zone, and the model was trained on it. The second issue was easily solved by dividing each weight by the sum of all weights. The predicted destination area is then chosen based on probability.

The results have shown a general improvement over the entire area. The major change is in the open water zone. The previous results could not be used due to the poor performance of the prediction. In [Bomberger et al., 2006], the accuracy was only 10 percent of the optimal predictor on eastward open sea prediction. It is now around 35 percent. The authors conclude that these results are good enough to justify trying the predictor in the middle of the ocean.

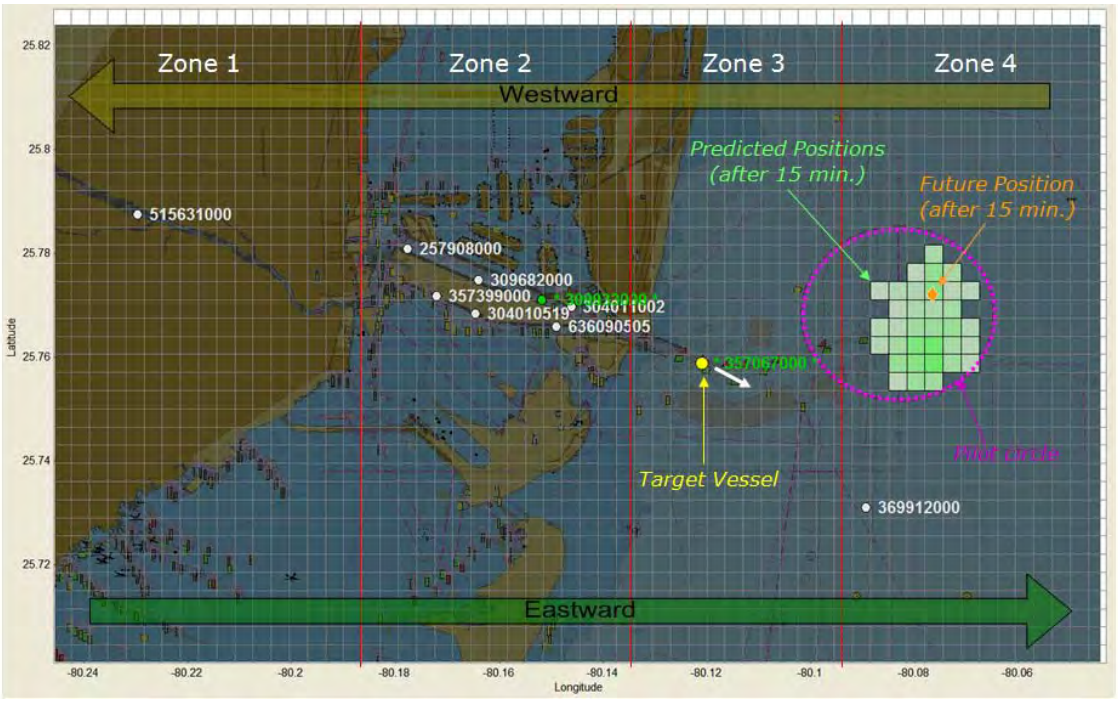

Figure 3 - Prediction of future vessel location [Bomberger et al., 2006]

3.3. Prototype systems

Some of the selected documents do not focus on a particular aspect of anomaly detection; instead, they present a global architecture. In this subsection, usable anomaly detection systems are presented. Some of them are already deployed or will be soon, and some are prototypes using real data. In all cases, they can be directly used by an operator to perform anomaly detection.

3.3.1. SeeCoast system

The prototype SeeCoast system is currently installed at the Joint Harbor Operations Center in Portsmouth, Virginia [Seibert, 2006]. It is built upon the Hawkeye system, which is deployed in more than six operational US Coast Guard Sector Command Centers. It fuses the video data with radar detections and AIS transponder data in order to generate composite fused tracks for vessels approaching the port, as well as for vessels already in the port. The desired end-state of the system is to automatically and reliably detect anomalies in the stream of maritime scene data while decreasing reliance on operator performance.

The system achieves this by adding a number of new capabilities, including the following: 1) video processing to detect, classify, and track vessels; 2) multi-sensor track correlation of video track data with the radar/AIS tracks; 3) automated camera control for track acquisition, ship size classification, and track maintenance; 4) more sophisticated rule-based track activity analysis to reduce operator workload; 5) learning-based track activity analysis to increase operator performance; and 6) display enhancements for improved situational awareness and forensic analysis.

In an attempt to adapt the system to the operator rather than adapting the operator to the system, an evolutionary approach was chosen: Hawkeye is improved instead of being replaced. Work is focused on new learning models and on developing more advanced video processing algorithms. The goal is to further improve the robustness of the system to local conditions and to reduce any remaining false alarms. At the time of the release of this article, there were still issues to address, but the SeeCoast system had passed its acceptance test plan.

3.3.2. SCANMARIS project

The SCANMARIS project is a technical component designed to detect abnormal vessel behaviours from an intelligent maritime traffic picture [Morel et al., 2009]. The traffic picture is processed from shoreline sensor data, deployed platform data and other available multisource information. The prototype is tested at the "Centre Régional Opérationnel de Surveillance et Sauvetage Corsen" (northwestern tip of France) on Ouessant traffic management.

The maritime traffic picture is accessed by both a rule engine and a learning engine. The rule engine processes the direct result of the data fusion process and analyzes it using defined rules established to detect abnormal behaviours. The learning engine, built on the AMAS (Adaptive Multi Agent System) theory, computes what combinations of these anomalies constitute relevant alarms. The system then issues these alarms to the human/machine interface. The operator manages each alarm and provides feedback to the learning engine.

There is also a threat-evaluation module called TAMARIS that quantifies the threat level of each alarm raised by the system. It is a supporting tool and uses multi-user, multi-pointer tactile interface devices to help collaborative groups of experts analyze and understand suspicious events. It uses a multi-hypothesis decision tree based on an ontology methodology and knowledge models from past experiences.

No performance assessment or lessons learned are provided. It is claimed that a system with such valued-added components would provide the existing maritime surveillance system with the means to efficiently combat criminal and illicit activities, as well as rule violations at sea.

3.3.3. LEPER

The work described in [Griffin, 2009-A] was supported by the Office of Naval Research (ONR) and was tested successfully at the Joint Interagency Task Force South (JIATF South). The system is called LEarning and Prediction for Enhanced Readiness (LEPER). It is built on the work of Christopher Griffin at the Applied Research Laboratory in Pennsylvania State University.

The system decomposes ship's trajectories into sequences of discrete squares from a military grid reference system. Using these sequences, the transition probabilities between grid locations are computed using HMM. To detect anomalies, the system predicts the position of vessels using a continuous model (i.e., speed and orientation) and compares it to the HMM prediction. If the distance between these two positions is above a predefined threshold, an anomaly is raised. However, the system does not take the shoreline into account, and the predicted position can be on land. Predicted trajectories are displayed with OpenMap or GoogleEarth.

LEPER is a research product and is not meant to be deployed. It has been tested successfully on limited open and classified maritime data: i.e., 100% of anomalies were detected. Future iterations will include multi-modal data, an SOA capability and multi-grid scales. It is claimed that the LEPER system can also include weather data.

3.3.4. SECMAR

Thales Underwater System leads a port security project called SECMAR (SECurity system to protect people, goods and facilities located in a critical MARitime area) [Géhant et al., 2009]. The main objective of the SECMAR project is to provide an awareness picture for the "close sea-side" areas in order to facilitate the task of the harbour surveillance and intervention teams. A prototype has to be designed for the strategic harbour of Fos-sur-Mer (Marseille, France).

With the help of subject matter experts, different kinds of threat scenarios are created. The designed behaviour can be either instantaneous (it depends only on the current instant) or temporal (the analysis requires many instants). Each scenario is expressed through an Esterel program, which compiles into a finite-state machine. The instantaneous observations are Esterel's program input signals, and the alarms are given by Esterel's program output signals.

The project consists of two phases. The first one is devoted to studies and the second to creating the prototype. The first phase is completed, and a complete maritime surveillance system has to be developed and integrated in the Fos-sur-Mer Gulf during Phase 2 of the SECMAR project. This project is scheduled to end in 2010 and results should follow.

3.3.5. FastC^{2}AP

This proof-of-concept system is built upon the DARPA Control of Agent Based Systems (CoABS) program [DARPA, 2005-G]. CoABS developed a technology with dynamic discovery and connection of military systems in heterogeneous environments by establishing a framework for integrating diverse legacy systems. This project is supported by many major organizations, such as DARPA and ONR.

The development of this system is focused on stimulating the adoption of agent-based technologies by the military services. To achieve this, the system makes extensive use of web services and collaboration tools. The "FastC^{2}AP Grid" middle-ware is where agents must be registered and where collaborators look for services. There are also state-of-the-art tools to configure single agents, or workflows for a series of agents. Here are some typical agents provided by FastC^{2}AP: Vessel Proximity Search, Abnormal Vessel Speed and Vessel Rendezvous Search [Bergeron, 2009; PSEG, 2007].

The goal of FastC^{2}AP is to establish a Common Maritime Operational Picture so that a single watchstander can manage more than 100 high-interest vessels with less than 1 percent false alarms in real-time identifications of anomalous behaviours. No information is given about how the system is intended to achieve this, what kind of technology is behind those agents or where the agents come from.

3.3.6. MALEF

The multi-agent learning framework (MALEF) presented in [Tozicka et al., 2008] is a generic, abstract framework used for distributed machine learning and data mining. This framework represents an attempt to capture complex forms of interaction between heterogeneous and/or self-interested learners. It can be used as a foundation for implementing complex interaction systems and reasoning mechanisms. This architecture allows agents to improve their knowledge using information provided by other learners in the system.

MALEF has been tested using data from the maritime domain (from the Frisia region) to detect anomalous vessel activities. This test assigns an agent for each geospatial region. These agents receive data from local vessels and perform data mining and machine learning tasks to build a model of normalcy. Agents can communicate models, data and hypotheses with their neighbourhood to achieve better explanation of the vessel behaviours. The anomaly detection focused on harbour visiting patterns and false vessel description.

It is claimed that this framework speeds up the building of a normalcy model because of the reduced quantity of data handled by each agent. Moreover, it does not seem to impact the performance of detection when sharing between agents is carefully selected. No performance result is provided in the paper.

3.4. Other research

Unlike the previous work presented so far, the following research efforts cannot be categorized or clustered easily. But although they may be just isolated efforts in the maritime domain, they were still generated by the anomaly detection community.

3.4.1. PANDA

The PANDA project is an initiative of DARPA. There is not much information on the status of this project. In fact, the only information available so far is the project information package. The mission of PANDA is to advance technologies and develop an architecture that will alert watchstanders to anomalous ship behaviour as it occurs, allowing them to detect potentially dangerous behaviour before it causes harm. The goal of PANDA is to automatically evaluate the behaviour of all larger-surface maritime vessels to determine which ones are deviating from their normal, expected behaviour in ways that may be indicative of an emerging threat.

PANDA will include research in motion-based pattern learning, prediction and activity monitoring, adaptive context, and anomaly processing and presentation. The project will be carried out in four phases. Phase I will focus on learning and detection, Phase II on automation, Phase III on integration, and Phase IV on technology scaling and transition.

3.4.2. Spatial characterization and spatio-semantic associations

The work presented in [Janeja et al., 2004] explains how to detect anomalous tracks using the characterization of regions surrounding locations of interest. The characterization of areas uses features like facilities, zoning, commerce, etc.

Feature vectors are used to identify anomalous shipping routes. The proposed model uses Voronoi diagrams to partition the areas into regions called "micro-neighbourhoods," each of which has a corresponding feature vector. These regions are defined by locations of interest like "city" or "port." Similarly-behaving micro-neighbourhoods are grouped to form "macro-neighbourhoods" and their feature vectors are averaged to form a composite feature vector. Then, it is possible to look for unexpected associations between a path and areas not on its expected trajectory, which may indicate deviations such as a stop-over. In addition, it is claimed that layers such as "drug zones" can be used to search for associations that can't be found by the traditional methods of spatial autocorrelation.

One example of application is given using ground shipping data. There is no performance evaluation or example on how to use this method in the maritime domain. However, it is mentioned that it could be used for the detection of anomalous cargo transshipment.

3.4.3. Maritime anomaly detection through interactive visualization

To improve the operator's confidence in a system, an anomaly detection process where the user is involved is proposed in [Riveiro et al., 2008]. Feedback and direction from the user are required for the whole process, from normalcy model building to detection refinement. There is a strong emphasis on the use of visual analytics.

Input data are first clustered using a self-organized map (SOM) that transforms multi-dimensional data vectors into two-dimensional clusters. However, it does not provide a complete solution to the anomaly detection problem, since there are many events that do not clearly fall into these well-defined clusters. Therefore, a Gaussian mixture model is used on top of the SOM. Once the model is built, the user adjusts it to reduce the number of false positives for each new anomaly detected.

Tests using synthetic data are provided. However, to make it possible to detect anomalies, the detection threshold had to be lowered to a level where false positives were too numerous. A simple solution is proposed to overcome this problem but no result is provided.

3.4.4. Semi-automatic ontology extension in the maritime domain

The work presented in [de Vries et al., 2008] addresses the problem of characterization of vessel type. Some ontologies are too high-level to divide ship classes into appropriate subclasses. The problem of heterogeneity prevents systems from deriving characterizations of vessel behaviour. To solve this problem, the authors propose extending the ontology in a semi-automatic way.

The first step in the proposed method is to characterize all AIS tracks by a hidden Markov model. Next, these models are clustered to form classes of ships behaving in the same way. For each

class, descriptions of ships are retrieved using their MMSI number. Using the most relevant description, major search engines are queried to retrieve potential hierarchy candidates to add to the ontology. These candidates are presented to a subject matter expert who will select the appropriate ones to be added.

No anomaly detection is presented in this paper. However, the extension of ontologies to characterize vessels can help greatly in putting the maritime domain in context. The validation of the process by a subject matter expert prevents potential errors. One example is given in which the vessel classification and the proposed hierarchy of vessel class are relevant.

3.4.5. An information theoretic approach to anomaly detection

[Barbará, 2009] presents an anomaly detection method based on information theory. The reason for this is to be able to combine multimodal data, while being robust to missing information, and to use a unique detection scheme for all types of anomalies. The entire process is data-driven and unsupervised, and no assumption is made about the distribution of data.

Data from vessels are considered as a multi-dimensional feature vector. Each feature is evaluated as a probability. Parzen windows are used for continuous values and counting occurrence for discrete ones. They calculate the strangeness of the feature vector by producing the sum of the surprise (i.e., the logarithm of the probability) of each component. Using a Parzen window, a distribution of the strangeness is approximated. An anomaly will be detected if the strangeness of a vector falls in the lower density of the distribution.

One of the major drawbacks of this method is that it requires relevant strangeness distributions. For example, fishing zones and sea lanes do not have the same kinds of vessels passing through them. For the anomaly detection to be accurate, different strangeness distributions are required. Some examples with real data are provided but there is no information on the performance of this method.

3.4.6. Dynamic network analysis for the detection of anomalies

[Carley et al., 2009] focus on the use of network analysis to increase maritime situational awareness; anomaly detection is only one of many topics addressed. The data exploited in this work include AIS reports, boarding reports, port entry/docking and land data.

The overall approach is to define a set of meta-networks and use them to identify critical entities using dynamic network analysis metrics. Anomaly detection is done on sets of ports that tended to be visited by the same ships and sets of ships that visited the same ports. Some of the methods described in Carley et al. [2009] are to identify critical owners, crews, passengers, ports and locations. Network analysis tools like network centrality, betweenness centrality and eigenvector centrality are used for this purpose.

No results are presented, other than two screen shots of reports. However, it is claimed that data provided by the Office of Naval Research were used. There is no explanation of how anomaly detection is performed or how it uses network analysis tools.

3.4.7. Compression and clustering for ship trajectory modelling

The method in [de Vries, van Someren, 2009] is an unsupervised way to construct a concise and effective model of a ship trajectory. The main goal of this approach is to model how ships are moving in a certain scenario or region, not to make an accurate kinematic model of ships. The model is constructed using vessel tracks taken from AIS data.

Tracks are first compressed using the Douglas-Peucker algorithm. This method comes from image processing and was invented to compress two-dimensional lines. However, it can also be used to compress time-series. In this application, tracks are compressed and then split into segments. After that, Affinity Propagation clustering (based on Euclidian distance of segment ends) is applied to create classes. To predict a future ship position, the last compressed segment of a track is matched to the closest class from the clustering process.

This method was tested on real AIS data, and evaluations were done for different kinds of ships. On average, it was superior to two baseline predictors for predictions done up to three hours in advance. With these performances, it could be possible to do anomaly detection based on the distance between predicted and actual positions of vessels.

3.4.8. Anomaly detection for maritime security

To handle the heavy and complex traffic in the port of Singapore, a bottom-up modelling approach is proposed in [Ma et al., 2009]. This method has been chosen to prevent subject matter experts from being diverted from their duty and to make it possible to automatically detect new unseen behaviours.

The first step in this method is to discretize the feature vector of each vessel contact. The speed, orientation and position will all be attributed a categorical value and be concatenate to form a feature word. Using a Hierarchical Dirichlet process clustering, classes of feature words, called activity classes, are created. The same clustering process is used on the sequence of feature vectors created from tracks, which are called behavioural models. A likelihood of occurrence is computed for each track using activity classes and behavioural model distribution. If the likelihood is too low (the threshold has to be tuned), an anomaly is raised.

From the experimental results, it is shown that this Dual Hierarchical Dirichlet process is able to model the complex maritime traffic of the Singapore Port. The resulting activity classes match closely with the knowledge base of the various locations and activities. In addition, with the behavioural models, it was possible to detect anomalous trajectories that correspond well to the subject matter expert's judgments.

3.4.9. Statistical analysis of motion patterns in AIS data

The application of statistical methods to detect anomalies and predict ship positions is presented in [Ristic et al., 2008]. This paper assumes that AIS data has already been processed and patterns have been extracted. It presents only the characterization of extracted patterns, i.e., the main sea lanes.

The statistical representation of sea lanes is built using an adaptive kernel density estimator (also known as the Parzen window method). The estimated probability distribution models the position and speed of vessels with a four-dimension vector by placing a kernel for each contact of a track. The anomaly detection process raises an anomaly if a new contact falls inside a low-density zone of the distribution (Fig. 4).

The paper also presents a method to predict future vessel positions. To do so, the method extracts the track that contributes the most to the distribution density at the current vessel position. A new distribution is created using samples of these tracks at the desired time in the future. This is claimed to be more accurate than a typical particle filter.

Figure 4 - Two-dimensional motion patterns of maritime traffic entering Port Jackson with anomaly decision boundaries [Ristic et al., 2008]

The paper presents preliminary results of ongoing research into behaviour analysis of vessels. No actual performance tests were provided, but the probability of false alarms of this detector can be evaluated numerically, thus providing its quantitative measure of performance. Also, the paper presents a fairly straightforward solution to motion prediction.

4. Assessment of anomaly detection methods

In this last part, anomaly detection methods from the previous part are assessed. Taxonomies of anomalies and detection methods are presented first. Next, a formal evaluation is done, followed by a mapping showing which methods address which anomalies. The last section provides comments and recommendations.

4.1. Taxonomies

Part 3 presented research papers related to anomaly detection. However, only twelve of them explicitly discuss anomaly detection methods. A distinction should be made between methods and techniques. Techniques are low-level tools for outlier detection. A review of the major family of techniques was presented earlier. Methods are procedures that use one or more techniques. Since these methods do not always have names, the titles of the papers in which they were presented will be used to identify them. Figure 5 shows a taxonomy of these methods grouped under the principal techniques they use.

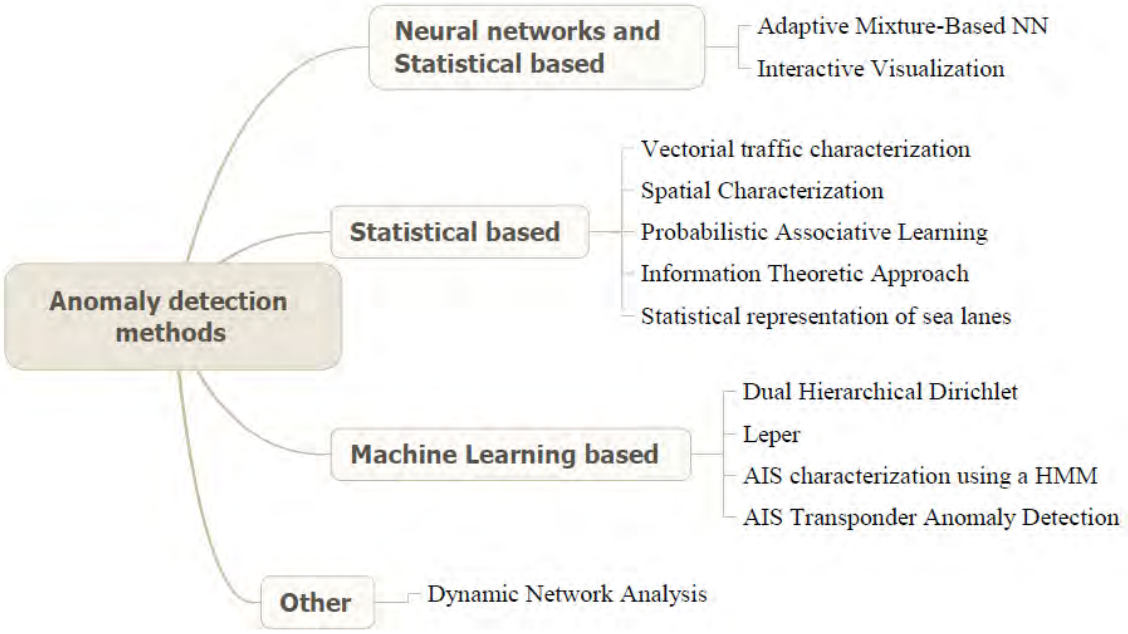

Figure 5–Anomaly detection methods taxonomy

These methods use different techniques and address different kinds of anomalies. Based on discussions with subject matter experts during workshops, [Davenport, 2008] identifies sixteen classes of kinematic anomalies. The taxonomy of these anomalies is shown in Figure 6. It encompasses all the anomalies detected by the methods presented here.

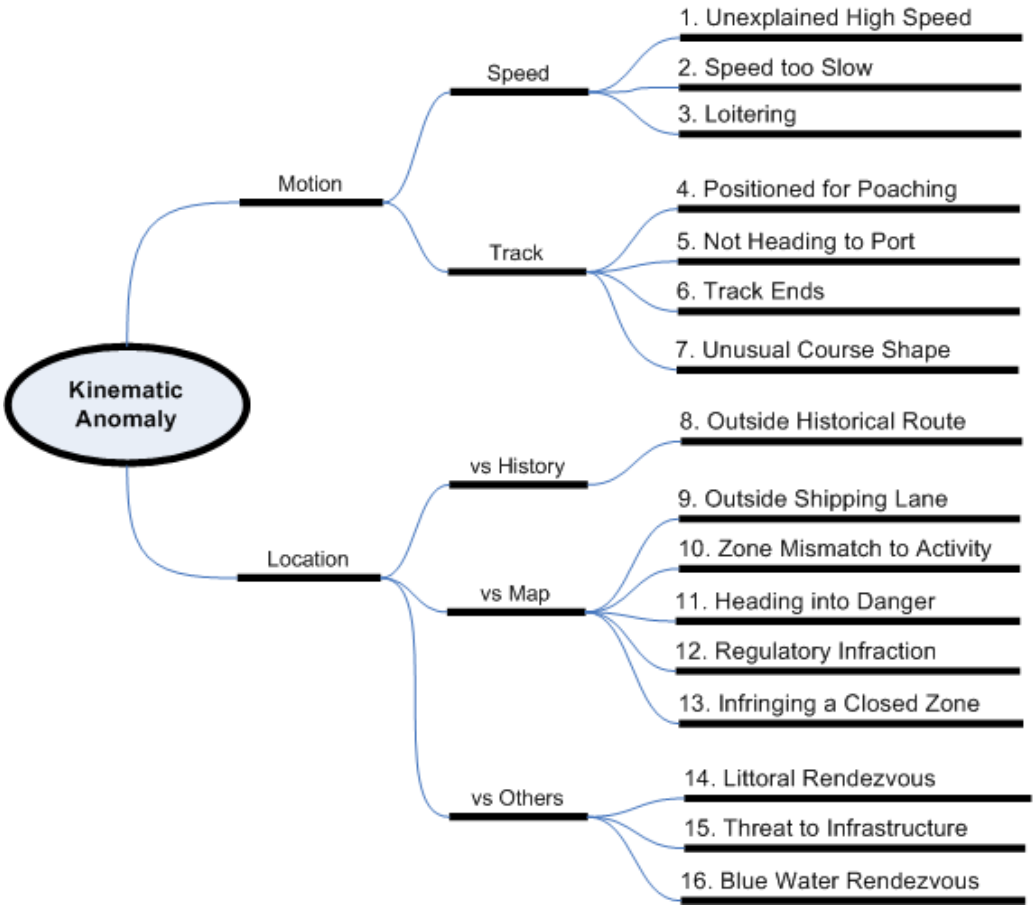

Figure 6 - Anomaly Taxonomy [Davenport, 2008]

It is now possible to create a mapping from these two taxonomies to show which methods address which anomalies. The choice was made to do so because these methods use many techniques and it is clear that it is meaningless to map anomalies against methods. The reason for this is that almost all problems in anomaly detection can be reduced to a point anomaly [Chandola et al., 2009] and be addressed by almost any technique. Moreover, this choice highlights combinations of techniques to solve anomalies. Table 1 presents this mapping.

Table 1 – Mapping of detected anomalies versus methods

	Dynamic Network Analysis	**Spatial Characterization**	**Adaptive Mixture-Based NN**	**Information Theoretic Approach**	**Dual Hierarchical Dirichlet**	**Interactive Visualization**	**Vectorial traffic characterization**	**Statistical representation of sea lanes**	**LEPER**	**Probabilistic Associative Learning**	**AIS characterization using a HMM**	**AIS Transponder Anomaly**
1. Unexplained high speed	●		●	●	●	●	●	●				
2. Speed too slow			●	●	●	●	●	●				
3. Loitering												
4. Positioned for poaching												
5. Not heading to port												
6. Track ends											●	●
7. Unusual course shape		●		●	●	●	●	●	●	●		
8. Outside historical route	●							●				
9. Outside shipping lane							●	●	●			
10. Zone mismatch to activity				●								
11. Heading into danger												
12. Regulatory infraction												
13. Infringing a Closed Zone		●										
14. Littoral rendezvous												
15. Threat to infrastructure												
16. Blue water rendezvous												

4.2. Methods evaluation

Since many methods address the same problems (see Table 1), an evaluation of the methods from the previous taxonomy is presented next. Using this evaluation, it will be possible to select a method that matches anomaly detection requirements. This assessment provides the following information about methods:

- What type of anomalies do they detect?
- What are their data requirements?
- What kind of techniques do they use?
- Expected success level
  - Unknown: not tested
  - High: claimed to be tested successfully
  - Medium: claimed to be tested; some issues remaining
  - Low: Major issues
- Comments

Table 2 – Dynamic Network Analysis assessment

Techniques used	Network analysis metrics
Anomalies detected	Outside historical route
Data requirements	List of locations/ports visited
Expected success level	Unknown
Comments	Assessment based on slide of a presentation; more investigation is required for a more complete assessment.

Table 3 – Vectorial traffic characterization assessment

Techniques used	Gaussian distribution Gaussian mixtures model Clustering (method not specified)
Anomalies detected	Unexplained high speed Speed too slow Unusual course shape Outside shipping lane
Data requirements	Ship position and speed (AIS)
Expected success level	High
Comments	Detect only anomalies on main sea lanes. Subject to high number of false positives.

Table 4 – Adaptive Mixture-Based NN assessment

Techniques used	Neural networks Gaussian mixtures model
Anomalies detected	Unexplained high speed Speed too slow
Data requirements	Ship position and speed (AIS)
Expected success level	Unknown

Comments	No evaluation of performance provided. Subject to high number of false positives. Anomaly detection is done on abstract feature vector and it can therefore be used for other anomalies if provided with different data.

Table 5 – Spatial Characterization assessment

Techniques used	Voronoi diagrams Jaccard coefficient
Anomalies detected	Zone mismatch to activity Infringing a closed zone
Data requirements	Geospatial location characterization (e.g., city with airport) Vessel position and cargo
Expected success level	Unknown
Comments	No evaluation of performance provided. Voronoi diagrams are based on location, so this method only applies to coastal region.

Table 6 – Information Theoretic Approach assessment

Techniques used	Information theoretic metric Parzen window Counting occurrence
Anomalies detected	Unexplained high speed Speed too slow Unusual course shape

Data requirements	Ship position and speed (AIS) Ship type distribution for geographical location
Expected success level	Medium
Comments	Anomaly detection is done on abstract feature vector and it can therefore be used for other anomalies if provided with different data. Subject to high number of false positives. This method works on a geographic grid and is subject to resolution problems.

Table 7 – Dual Hierarchical Dirichlet assessment

Techniques used	Hierarchical Dirichlet process Gibbs sampling
Anomalies detected	Unexplained high speed Speed too slow Unusual course shape
Data requirements	Ship position and speed
Expected success level	High
Comments	Anomaly detection is done on abstract feature vector so can be used for other anomalies if provided with different data. This method works on a geographic grid and is subject to resolution problems.

Table 8 – Statistical representation of sea lanes assessment

Techniques used	Parzen window
Anomalies detected	Unexplained high speed Speed too slow

	Unusual course shape Outside shipping lane
Data requirements	Ship position and speed
Expected success level	Unknown
Comments	Anomaly detection is done on abstract feature vector and it can therefore be used for other anomalies if provided with different data. No evaluation of performance provided. Subject to high number of false positives. Only detects anomalies on main sea lanes.

Table 9 – Interactive Visualization assessment

Techniques used	Gaussian mixture Neural networks
Anomalies detected	Unexplained high speed Speed too slow Unusual course shape
Data requirements	Ship position and speed
Expected success level	Medium
Comments	Subject to high number of false positives.

Table 10 – LEPER assessment

Techniques used	Hidden Markov model
Anomalies detected	Unusual course shape

Data requirements	Ship position and speed
Expected success level	Medium
Comments	This method works on a geographical grid and is subject to resolution problems.

Table 11 – Probabilistic Associative Learning assessment

Techniques used	Hebbian learning Counting occurrence
Anomalies detected	Unusual course shape
Data requirements	Ship position and speed
Expected success level	Medium
Comments	This method works on a geographical grid and is subject to resolution problems. Not tested in open sea.

Table 12 – AIS characterization using a HMM assessment

Techniques used	Hidden Markov model
Anomalies detected	Track ends
Data requirements	Ship position
Expected success level	Low
Comments	Work proven unusable by [Guerriero et al.]

Table 13 – AIS Transponder Anomaly assessment

Techniques used	Hidden Markov model Neyman-Pearson rule Maximum *a posteriori* (MAP) decision rule
Anomalies detected	Track ends
Data requirements	Ship position
Expected success level	Low
Comments	Work proven unusable by [Guerriero et al.]

4.3. Comments and recommendations

The list of anomaly detection methods presented here does not cover all the types of anomalies that are of interest in the maritime domain. Moreover, many methods try to solve the same problem with different techniques. The fact is that there is no single method for handling each anomaly. The problem of anomaly detection is that of finding a way to reduce the problem to a state where anomaly detection techniques can be applied and compared.

Anomaly detection is not a new topic; it has been studied it for years. Techniques are mature and well understood. Many libraries exist (some of which are free) that can be added to computer programs to perform anomaly detection. With the objective of covering all the anomalies presented earlier (see Figure 6), one should not focus on anomaly detection techniques per se, but more on the reduction of the problem to a more simple form.

Speed and position anomalies are often topics of research because the reduction problem is inexistent; people just compare the efficiency of different techniques. More research is needed to handle maritime anomalies not covered so far (see Table 1). For a particular anomaly, the challenge is to transform input data to a point anomaly problem where known techniques can easily detect outliers.

5. Conclusion

This document was written with the objective of presenting a high-level overview of the research in the field of maritime anomaly detection. The process of anomaly detection was depicted in a conceptual way in the first part, using material from conferences, workshops and scientific papers. The second part used a more formal approach to describe in a concrete way the methods that address issues highlighted in the first part. The selected literature review was structured around specific organizations known to be active in maritime anomaly detection, various MAD systems, and other relevant research activities. Lastly, an assessment of anomaly detection methods was presented, together with comments and recommendations, in order to identify gaps.

There is a great deal of research activity on the topic of anomaly detection. Many detection techniques are well understood and used in many domains. The maritime domain offers plenty of opportunities for research and innovation to develop or improve computer systems that will use these techniques. There is a need for MAD, and too many of the current efforts are focused on the same area, i.e., unusual speed or position. Future research work should try to fill the gaps in the coverage of the different anomaly types and give computer systems extended capabilities, such as threat assessment, that will be useful to operators.

References

[Baldacci, 2008-A], Baldacci, A., *AIS Emission Anomaly Detection in Support of Maritime Surveillance*, Technical Report, NURC-FR-2008-020, July 2008.

[Baldacci, 2008-B], Baldacci, A., *Automated Identification System Anomaly Simulator*, Memorandum Report, NURC-MR-2008-002, June 2008.

[Baldacci, 2008-C], Baldacci, A., *Anomaly Detection for MSA*, NATO TIDE SPRINT, Virginia Beach, USA, 27-31 October 2008.

[Baldacci, Fabiani, 2008], Baldacci, A., and Fabiani, A., *Contact-Based AIS Coverage Estimation and Distribution*, Technical Report, NURC-MR-2008-001, March 2008.

[Baldacci, Carthel, 2009], Baldacci, A., and Carthel, C., *Maritime traffic characterization with the Automated Identification System*, Technical Report, NURC-FR-2009-008, May 2009.

[Baldacci et al., 2009], Baldacci, A., Cappelletti, M., Carthel, C., and Coraluppi, S., *AIS Transponder Anomaly Detection for Maritime Situational Awareness*, NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness (NATO MSA 2009), NATO Undersea Research Centre (NURC), La Spezia, Italy, 15-17 September 2009.

[Barbará, 2009], Barbará, D., *An Information Theoretic Approach to Anomaly Detection*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Barrett, 2009], Barrett, W., *Evaluating Behavioral Indicators for Maritime Domain Awareness*, Presentation to TTCP MAR AG-8, 06 April, 2009.

[Bergeron, 2009], Bergeron, V., FASTC^{2}AP Agents Analysis, April 2009.

[Bomberger et al., 2006], Bomberger, N.A., Rhodes, B.J., Seibert, M., and Waxman, A.M., *Associative Learning of Vessel Motion Patterns for Maritime Situation Awareness*, In Proceedings of the 9th International Conference on Information Fusion (Fusion 2006), Florence, Italy, July 10-13, 2006.

[Boner, 2009], Boner, C., *Cargo Domain Technologies and Technology Gaps*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Boraz, 2009], Boraz, Cdr S., *Program Executive Office, Command, Control, Communications, Computers and Intelligence (PEO C4I) – Maritime Domain Awareness*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Carley et al., 2009], Carley, K.M., Davis, G.B., and Olson, J., *Dynamic Network Analysis for the Detection of Anomalies to Support Maritime Analysis*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Chandola et al., 2009], Chandola V., Banerjee A., and Kumar, V., Anomaly detection: A survey, ACM Computing Surveys (CSUR), v.41 n.3, pp. 1-58, July 2009.

[DARPA, 2005-A], DARPA, *Proposal Information Package (PIP) - Predictive Analysis for Naval Deployment Activities (PANDA) - BAA 05-44*, 28 September 2005.

[DARPA, 2005-B], DARPA, *Frequently Asked Questions (FAQ) - Predictive Analysis for Naval Deployment Activities (PANDA) - BAA 05-44*, 12 October 2005.

[DARPA, 2005-C], DARPA, Predictive Analysis for Naval Deployment Activities (PANDA) - Mission, 2005.

[DARPA, 2005-D], DARPA, Predictive Analysis for Naval Deployment Activities (PANDA) - Vision, 2005.

[DARPA, 2005-E], DARPA, Predictive Analysis for Naval Deployment Activities (PANDA) - Goals, 2005.

[DARPA, 2005-F], DARPA, Predictive Analysis for Naval Deployment Activities (PANDA) - Challenges, 2005.

[DARPA, 2005-G], DARPA, *Fast Connectivity for Coalitions and Agents Project*, Fact Sheet, 2005.

[Davenport, 2008], Davenport, M., *Kinematic Behaviour Anomaly Detection (KBAD) - Final Report*, MacDonald Dettwiler and Associates Ltd, DRDC CORA CR 2008-002, DRDC CORA Project Manager: Neil Carson, April 2008.

[de Vries et al., 2008], de Vries, G., Malaisé, V., van Someren, M., Adriaans, P., and Schreiber, G., *Semi-Automatic Ontology Extension in the Maritime Domain*, The 20th Belgian-Netherlands Conference on Artificial Intelligence (BNAIC 2008), University of Twente, Enschede, the Netherlands, October 30-31, 2008.

[de Vries, van Someren, 2009], de Vries, G., and van Someren, M., *Unsupervised Ship Trajectory Modeling and Prediction Using Compression and Clustering*, Proceedings of the 18th Annual Belgian-Dutch Conference on Machine Learning, pages 7-12, Tilburg, May 2009.

[Garagic et al., 2009], Garagic, D., Rhodes, B.J., Bomberger, N.A., and Zandipour, M., *Adaptive Mixture-Based Neural Network Approach for Higher-Level Fusion and Automated Behavior Monitoring*, NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness (NATO MSA 2009), NATO Undersea Research Centre (NURC), La Spezia, Italy, 15-17 September 2009.

[Géhant et al., 2009], Géhant, M., Roy, V., Marmorat, J.-P., and Bordier, M., *A Behaviour Analysis Prototype for Application to Maritime Security*, NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness (NATO MSA 2009), NATO Undersea Research Centre (NURC), La Spezia, Italy, 15-17 September 2009.

[Griffin, 2009-A], Griffin, C., *Learning and Prediction for Enhanced Readiness: An ONR Office 31 Program*, Presentation to TTCP MAR AG-8, 06 April, 2009.

[Griffin, 2009-B], Griffin, C., *Multi-Layer Statistical Methods for Learning Maritime Behavior*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Griffin et al., 2009], Griffin, C., Clark, D., Deans, L., Fahlman, S., Nevell, D., and Roy, J., *Research Gaps and Taxonomy for Anomaly Detection with Emphasis on the Maritime Domain*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Guerriero et al.], Guerriero M., Coraluppi, S., and Carthel, C., *Analysis of AIS intermittency and vessel characterization using a Hidden Markov Model*, Technical Report, NURC-FR-2010-002, January 2010.

[Hodge et al., 2004], Hodge V., and Austin, J., A Survey of Outlier Detection Methodologies, Artificial Intelligence Review, v.22 n.2, pp. 85-126, October 2004

[Janeja et al., 2004], Janeja, V.P., Atluri, V., and Adam, N.R., *Detecting Anomalous Geospatial Trajectories through Spatial Characterization and Spatio-Semantic Associations*, Proceedings of the 2004 annual conference on digital government research, 2004.

[Kessler, 2009], Kessler, O., *Workshop on Detection of Anomalous Behaviors in Maritime Environments*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Ma et al., 2009], Ma, K.-T., Ng, G.-W., Wang, X., and Grimson, W.E.L., *Anomaly detection for Maritime Security*, NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness (NATO MSA 2009), NATO Undersea Research Centre (NURC), La Spezia, Italy, 15-17 September 2009.

[Moore, 2005], Moore, K.E., *BAA 05-44 - Predictive Analysis for Naval Deployment Activities (PANDA) - Briefing to Industry: PANDA Overview*, 16 September 2005.

[Morel et al., 2009], Morel, M., George, J.-P., Jangal, F., Napoli, A., Giraud, M.-A., Botalla, M., and Littaye, A., *SCANMARIS Project – Detection of Abnormal Vessel Behaviours*, NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness (NATO MSA 2009), NATO Undersea Research Centre (NURC), La Spezia, Italy, 15-17 September 2009.

[PSEG, 2007], Pacific Science & Engineering Group, *FastC^{2}AP Training: FastC^{2}AP Portal Operator Training*, Fast Connectivity for Coalitions and Agents Project (FastC^{2}AP) - Build 2.4, Training Developed by Pacific Science & Engineering Group, 2007.

[Ristic et al., 2008], Ristic, B., La Scala, B., Morelande, M., and Gordon, N., *Statistical Analysis of Motion Patterns in AIS Data - Anomaly Detection and Motion Prediction*, in Proceedings of The 11th International Conference on Information Fusion (Fusion 2008), Cologne, Germany, June 30 – July 03, 2008.

[Riveiro et al., 2008], Riveiro, M., Falkman, G., and Ziemke, T., *Improving Maritime Anomaly Detection and Situation Awareness Through Interactive Visualization*, Proceedings of the 11th International Conference on Information Fusion (Fusion 2008), Cologne, Germany, 30 June - 03 July, 2008.

[Rhodes, 2007], Rhodes, B.J., *Biologically-Inspired Approaches to Higher-Level Information Fusion*, Proceedings of the 10th International Conference on Information Fusion (Fusion 2007), Quebec, Canada, 9-12 July 2007.

[Rhodes et al., 2005], Rhodes, B.J., Bomberger, N.A., Seibert, M., and Waxman, A.M., *Maritime Situation Monitoring and Awareness Using Learning Mechanisms*, In Proceedings of IEEE MILCOM 2005 Military Communications Conference, Atlantic City, NJ, USA, October 17-20, 2005.

[Rhodes et al., 2007], Rhodes, B.J., Bomberger, N.A., and Zandipour, M., *Probabilistic Associative Learning of Vessel Motion Patterns at Multiple Spatial Scales for Maritime Situation Awareness*, Proceedings of the 10th International Conference on Information Fusion (Fusion 2007), Quebec, Canada, 9-12 July 2007.

[Seibert, 2009], Seibert, M., *Maritime Anomaly Detection*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Seibert, 2006] Seibert, M., Rhodes, B.J., Bomberger, N.A, Beane, P.O., Sroka, J.J., Kogel, W., Kreamer, W., Stauffer, C., Kirschner, L., Chalom, E., Bosse, M., and Tillson, R., *SeeCoast port surveillance*, *Proceedings of SPIE Vol. 6204: Photonics for Port and Harbor Security II* Orlando, FL, USA, 18–19 April, 2006.

[Sisk et al., 2009], Sisk, B., Woessner, B., and Weng, Y., *Detecting Unusual Events in the Maritime Domain – Program Review*, Workshop on Detection of Anomalous Behaviors in Maritime Environments, Carnegie Mellon University, 25-26 June 2009.

[Smith et al., 2009], Smith, A.J.E., Anitori, L., Bergmans, J., Colin, M., van Iersel, M., Liem, K.D., Schwering, P.B.W., van Sweeden, R., and Vullings, H.J.L.M., *Overview of maritime situational awareness research at the Netherlands Organization for Applied Scientific Research TNO*, NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness (NATO MSA 2009), NATO Undersea Research Centre (NURC), La Spezia, Italy, 15-17 September 2009.

[Tarsus, 2009], Tarsus, D., *Criteria Taken Into Account at a Vessel Rendezvous*, MSA Analyst, CC MAR Naples, 04 February 2009.

[Tidepedia, 2009], Tidepedia, *Smart Agent Development – Ideas from the Scenario Group*, Ede, 2009.

[Tozicka et al., 2008], Tozicka, J., Rovatsos, M., Pechoucek, M., and Urban, S., *MALEF: Framework for Distributed Machine Learning and Data Mining*, International Journal of Intelligent Information and Database Systems. 2008, vol. 2, pp. 6-24.

[Walden, 2006], Walden, R., *Automated Situation Assessment in Maritime Combat Systems*, Presentation to TTCP MAR TP1, Adelaide, Australia, 6-10 November 2006.

List of symbols/abbreviations/acronyms/initialisms

AIS	Automatic Identification System
AMAS	Adaptive Multi Agent System
ARP	Applied Research Program
ARTMAP	Adaptive Resonance Theory Map
CANCOM	Canada Command
CCMAR	Center of Marine Sciences
CCI	Centre de commandement interarmées
CROI	Centres régionaux d'opérations interarmées
DAM	Détection d'anomalies maritimes
DARPA	Defense Advanced Research Projects Agency
DND	Department of National Defence
DRDC	Defence Research & Development Canada
DRDKIM	Director Research and Development Knowledge and Information Management
GEC	Gestion et d'Exploitation de la Connaissance
HMM	Hidden Markov Models
JIATF-S	Joint Interagency Task Force South.
JCC	Joint Command Centre
KME	knowledge management and exploitation
LNG	Liquefied Natural Gas
MAD	Maritime Anomaly Detection
MALEF	Multi-Agent Learning Framework
MAP	Maximum A Posteriori
MMSI	Maritime Mobile Service Identity
MS09	Maritime Surveillance 09
MSDS	Maritime Surveillance Data Simulator
NATO	North Atlantic Treaty Organization
ONR	Office of Naval Research
OGC	Open Geospatial Consortium
NURC	NATO Undersea Research Center

PANDA	Predictive Analysis for Naval deployment Activities
PRA	Programme de recherche appliquée
R&D	Research & Development
RJOCs	Regional Joint Operations Centres
SOA	Services-Oriented Architectures
SOM	Self-Organized Map
UAV	Unmanned Aerial Vehicles
WMD	Weapon of Mass Destruction

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Defence R&D Canada – Valcartier
2459 Pie-XI Blvd North
Quebec (Quebec)
G3J 1X5 Canada

2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.)

UNCLASSIFIED

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Maritime Anomaly Detection: Domain Introduction and Review of Selected Literature

4. AUTHORS (last name, followed by initials – ranks, titles, etc. not to be used)

Martineau, E.; Roy, J.

5. DATE OF PUBLICATION (Month and year of publication of document.)

October 2011

6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)

62

6b. NO. OF REFS (Total cited in document.)

50

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g., interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

Defence R&D Canada – Valcartier
2459 Pie-XI Blvd North
Quebec (Quebec)
G3J 1X5 Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)

9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DRDC Valcartier TM 2010-460

10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

Unlimited

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

Early in the conduct of Project 11hg (Collaborative Knowledge Exploitation for Maritime Domain Awareness) at Defence R&D Canada, anomaly detection in the maritime domain has been identified by the operators/analysts of the operational community as an important aspect requiring research and development. A number of R&D activities have thus been undertaken under the project to specifically investigate maritime anomaly detection (MAD). This Technical Memorandum reports on one of these activities. It first provides a high-level introduction to the domain, and then presents a review of selected literature on the subject. Different views of the field are presented, starting with a description of the various steps of MAD, followed by a discussion of four interrelated goals of MAD. Current gaps in MAD are identified from the data and information, processing and systems perspectives. The selected literature review is structured around specific organizations known to be active in maritime anomaly detection, various MAD systems, and other relevant research activities. A high-level assessment of the methods for MAD that were found in the reviewed literature completes the discussion.

Tôt dans l'exécution du projet 11hg (Collaborative Knowledge Exploitation for Maritime Domain Awareness) à la R&D pour la Défense du Canada, la détection d'anomalies dans le domaine maritime fut identifiée par les opérateurs/analystes de la communauté opérationnelle comme un aspect important qui nécessite de la recherche et du développement. Plusieurs activités de R&D furent entreprises dans le cadre du projet pour investiguer spécifiquement la détection d'anomalies maritimes (DAM). Ce mémorandum technique fait état de l'une de ces activités. Il fournit d'abord une introduction au domaine et ensuite présente une revue de littérature sélectionnée sur le sujet. Différentes visions du domaine sont présentées en commençant par la description des différentes étapes de la DAM suivie par une discussion sur ses quatre objectifs. Les lacunes actuelles de la DAM sont identifiées dans la perspective des données et de l'information, du traitement et des systèmes. La revue de littérature est structurée autour d'organisations spécifiques connues pour être actives dans la DAM, les divers systèmes de gestion de la DAM et autres activités de recherche pertinentes. Une évaluation de haut niveau des méthodes de la DAM qui furent identifiées dans la revue de littérature complète la discussion.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

maritime anomaly detection, maritime domain awareness

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

www.drdc-rddc.gc.ca